

AWS-Certified-Solutions-Architect-Professional Dumps

Amazon AWS Certified Solutions Architect Professional

<https://www.certleader.com/AWS-Certified-Solutions-Architect-Professional-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2, an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand Instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:

- Address scalability issues.
- Increase the level of concurrency.
- Eliminate licensing costs.
- Improve reliability.

Which set of steps should the solutions architect take?

- A. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- B. Convert the Oracle database into a single Amazon RDS reserved DB instance.
- C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- D. Create two additional copies of the database instance, then distribute the databases in separate AZs.
- E. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- F. Convert the tables in the Oracle database into Amazon DynamoDB tables.
- G. Convert the On-Demand Instances into Spot Instances to reduce costs for the front end
- H. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Answer: C

Explanation:

Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 2

- (Exam Topic 1)

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region. What should a solutions architect do to meet these requirements?

- A. Create a new developer account
- B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
- C. Enforce a tagging policy that denotes Region affinity.
- D. Create an SCP that denies the launch of all EC2 instances except t3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
- E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
- F. Create an IAM policy that allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 3

- (Exam Topic 1)

A company wants to control its cost of Amazon Athena usage. The company has allocated a specific monthly budget for Athena usage. A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount. Which solution will meet these requirements?

- A. Use AWS Budget
- B. Create an alarm (or when the cost of Athena usage reaches the budgeted amount for the month)
- C. Configure AWS Budgets actions to deactivate Athena until the end of the month.
- D. Use Cost Explorer to create an alert for when the cost of Athena usage reaches the budgeted amount for the month
- E. Configure Cost Explorer to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Use AWS Trusted Advisor to track the cost of Athena usage
- G. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to deactivate Athena until the end of the month whenever the cost reaches the budgeted amount for the month
- H. Use Athena workgroups to set a limit on the amount of data that can be scanned
- I. Set a limit that is appropriate for the monthly budget and the current pricing for Athena.

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS CloudFormation. The applications are in Typescript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts. Developers from the newly acquired company are hesitant to move their applications under CloudFormation because it would require them to learn a new domain-specific language and eliminate their access to language features, such as looping. How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

- A. Create CloudFormation templates and re-use parts of the Python scripts as instance user data
- B. Use the AWS Cloud Development Kit (AWS CDK) to deploy the application using these templates
- C. Incorporate the AWS CDK into CodePipeline and deploy the application to AWS using these templates.
- D. Use a third-party resource provisioning engine inside AWS CodeBuild to standardize the deployment processes of the existing and acquired companies
- E. Orchestrate the CodeBuild job using CodePipeline.
- F. Standardize on AWS OpsWorks

- G. Integrate OpsWorks with CodePipelin
- H. Have the developers create Chef recipes to deploy their applications on AWS.
- I. Define the AWS resources using Typescript or Pytho
- J. Use the AWS Cloud Development Kit (AWS CDK) to create CloudFormation templates from the developers' code, and use the AWS CDK to create CloudFormation stack
- K. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

A solution architect needs to deploy an application on a fleet of Amazon EC2 instances. The EC2 instances run in private subnets in An Auto Scaling group. The application is expected to generate logs at a rate of 100 MB each second on each of the EC2 instances.

The logs must be stored in an Amazon S3 bucket so that an Amazon EMR cluster can consume them for further processing The logs must be quickly accessible for the first 90 days and should be retrievable within 48 hours thereafter.

What is the MOST cost-effective solution that meets these requirements?

- A. Set up an S3 copy job to write logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a NAT instance within the private subnets to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier.
- B. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive
- C. Set up an S3 batch operation to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a NAT gateway with the private subnets to connect to Amazon S3 Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive
- D. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier

Answer: C

NEW QUESTION 6

- (Exam Topic 1)

A solutions architect at a largo company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization m AWS Organizations The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the interne) The company deploys resources only Into a single AWS Region The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone Which solution meets these requirements?

- A. Creates a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Create an Auto Scaling group of Amazon EC2 Instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region Modify all default routes to point to the proxy's Auto Scaling group
- B. Create a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Use an AWS Network Firewall firewall for rule-based filtering Create Network Firewall endpoints In each Availability Zone Modify all default routes to point to the Network Firewall endpoints
- C. Create an AWS Network Firewall firewal for rule-based filtering in each AWS account Modify all default routes to point to the Network Firewall firewalls in each account.
- D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/> <https://aws.amazon.com/blogs/networking-and-content-delivery/deploy-centralized-traffic-filtering-using-aws-n>

NEW QUESTION 7

- (Exam Topic 1)

A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:

- The database must use strong, randomly generated passwords stored in a secure AWS managed service.
- The application resources must be deployed through AWS CloudFormation.
- The application must rotate credentials for the database every 90 days.

A solutions architect will generate a CloudFormation template to deploy the application.

Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

- A. Generate the database password as a secret resource using AWS Secrets Manage
- B. Create an AWS Lambda function resource to rotate the database passwor
- C. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.
- D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Stor
- E. Create an AWS Lambda function resource to rotate the database passwor
- F. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.
- G. Generate the database password as a secret resource using AWS Secrets Manage
- H. Create an AWS Lambda function resource to rotate the database passwor
- I. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.
- J. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Stor
- K. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/how-to-securely-provide-database-credentials-to-lambda-functions-by-us>

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html

NEW QUESTION 8

- (Exam Topic 1)

A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.

What changes to the current architecture will reduce operational overhead and support the product release?

- A. Create an EC2 Auto Scaling group behind an Application Load Balance
- B. Create additional read replicas for the DB instance
- C. Create Amazon Kinesis data streams and configure the application services to use the data stream
- D. Store and serve static content directly from Amazon S3.
- E. Create an EC2 Auto Scaling group behind an Application Load Balance
- F. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling
- G. Create Amazon Kinesis data streams and configure the application services to use the data stream
- H. Store and serve static content directly from Amazon S3.
- I. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balance
- J. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling
- K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster
- L. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- M. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balance
- N. Create additional read replicas for the DB instance
- O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster
- P. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:

Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

NEW QUESTION 9

- (Exam Topic 1)

A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10,000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.

Which design should a solutions architect implement?

- A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Log
- B. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processed
- C. Place processed images in another directory in Amazon EFS
- D. Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.
- E. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instance
- F. Enable Amazon CloudFront and configure the origin to be the ALB in front of the EC2 instances.
- G. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SQS) queue
- H. Create a fleet of Amazon EC2 instances to pull messages from the SQS queue to process the images and place them in another S3 bucket
- I. Use Amazon CloudWatch metrics for queue depth to scale out EC2 instance
- J. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.
- K. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume mounted to a fleet of Amazon EC2 Spot instance
- L. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processed
- M. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to scale out EC2 instance
- N. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.

Answer: C

NEW QUESTION 10

- (Exam Topic 1)

A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC, and is encrypted with an AWS KMS key. A solutions architect has verified that the developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid.

Which additional step should the solutions architect take to troubleshoot this issue?

- A. Ensure that blocking all public access has not been enabled in the S3 bucket.
- B. Verify that the IAM role has permission to decrypt the referenced KMS key.
- C. Verify that the IAM role has the correct trust relationship configured.
- D. Check that local firewall rules are not preventing access to the S3 endpoint.

Answer: B

NEW QUESTION 10

- (Exam Topic 1)

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalogue

page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times. Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

- A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
- B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionalit
- C. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- D. Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route 53 health check against the product page to evaluate full application functionalit
- E. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- F. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
- G. Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Answer: BE

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/health-checks-types.html>

NEW QUESTION 15

- (Exam Topic 1)

A company has an application that generates reports and stores them in an Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.

Which set of actions will immediately remediate the security issue without impacting the application's normal workflow?

- A. Create an AWS Lambda function that applies a deny all policy for users who are not authenticated. Create a scheduled event to invoke the Lambda function.
- B. Review the AWS Trusted Advisor bucket permissions check and implement the recommended actions.
- C. Run a script that puts a private ACL on all of the objects in the bucket.
- D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcls option to TRUE on the bucket.

Answer: D

Explanation:

The S3 bucket is allowing public access and this must be immediately disabled. Setting the IgnorePublicAcls option to TRUE causes Amazon S3 to ignore all public ACLs on a bucket and any objects that it contains. The other settings you can configure with the Block Public Access Feature are:

- o BlockPublicAcls – PUT bucket ACL and PUT objects requests are blocked if granting public access.
- o BlockPublicPolicy – Rejects requests to PUT a bucket policy if granting public access.
- o RestrictPublicBuckets – Restricts access to principles in the bucket owners' AWS account. <https://aws.amazon.com/s3/features/block-public-access/>

NEW QUESTION 16

- (Exam Topic 1)

A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.

To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.

Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instance
- B. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- C. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attache
- D. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- E. Create separate route tables for production and development traffi
- F. Delete each account's association and route propagation to the default AWS Transit Gateway route tabl
- G. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
- H. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attache
- I. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/tgw/vpc-tgw.pdf>

NEW QUESTION 18

- (Exam Topic 1)

A company wants to migrate a 30 TB Oracle data warehouse from on premises to Amazon Redshift The company used the AWS Schema Conversion Tool (AWS SCT) to convert the schema of the existing data warehouse to an Amazon Redshift schema The company also used a migration assessment report to identify manual tasks to complete.

The company needs to migrate the data to the new Amazon Redshift cluster during an upcoming data freeze period of 2 weeks The only network connection between the on-premises data warehouse and AWS is a 50 Mops internet connection

Which migration strategy meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instance
- B. Authorize the public IP address of the replication instance to reach the data warehouse through the corporate firewall Create a migration task to run at the beginning of the data freeze period.
- C. Install the AWS SCT extraction agents on the on-premises server
- D. Define the extract, upload, and copy tasks to send the data to an Amazon S3 bucket
- E. Copy the data into the Amazon Redshift cluster
- F. Run the tasks at the beginning of the data freeze period.
- G. Install the AWS SCT extraction agents on the on-premises server
- H. Create a Site-to-Site VPN connection Create an AWS Database Migration Service (AWS DMS) replication instance that is the appropriate size Authorize the IP address of the replication instance to be able to access the on-premises data warehouse through the VPN connection
- I. Create a job in AWS Snowball Edge to import data into Amazon S3 Install AWS SCT extraction agents on the on-premises servers Define the local and AWS Database Migration Service (AWS DMS) tasks to send the data to the Snowball Edge device When the Snowball Edge device is returned to AWS and the data is available in Amazon S3, run the AWS DMS subtask to copy the data to Amazon Redshift.

Answer: D

Explanation:

AWS Database Migration Service (AWS DMS) can use Snowball Edge and Amazon S3 to migrate large databases more quickly than by other methods

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_LargeDBs.html

https://www.calctool.org/CALC/prof/computing/transfer_time

NEW QUESTION 22

- (Exam Topic 1)

A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.

The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.

Which solution meets these requirements?

- A. Configure scan on push on the repository
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity findings. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
- C. Configure scan on push on the repository. Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Lambda function when a new message is added to the SQS queue. Use the Lambda function to delete the image tag for images that have Critical or High severity findings.
- D. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- E. Schedule an AWS Lambda function to start a manual image scan every hour. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke another Lambda function when a scan is complete.
- F. Use the second Lambda function to delete the image tag for images that have Critical or High severity findings.
- G. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- H. Configure periodic image scan on the repository. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity findings.
- I. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: C

NEW QUESTION 23

- (Exam Topic 1)

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue
- E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- F. Modify the application to use Amazon DynamoDB instead of Amazon RDS
- G. Configure Auto Scaling for the DynamoDB table
- H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization
- I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- J. Update the application to use a Redis task queue instead of the in-memory queue
- K. Build a Docker container image for the application
- L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis
- M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:

The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the defacto answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

NEW QUESTION 28

- (Exam Topic 1)

A company has 50 AWS accounts that are members of an organization in AWS Organizations. Each account contains multiple VPCs. The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account. Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.

Which combination of steps will meet these requirements? (Select TWO)

- A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
- B. From the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
- C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account
- D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
- E. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a peering transit gateway attachment in a member account
- F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
- G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

Explanation:

<https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/>

NEW QUESTION 30

- (Exam Topic 1)

A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments.

Which combination of steps will meet these requirements? (Select THREE).

- A. Mirror the application code to an AWS CodeCommit Git repository
- B. Use the repository to build EC2 AMIs.
- C. Produce multiple EC2 AMI
- D. one for each environment, for each release.
- E. Produce one EC2 AMI for each release for use across all environments.
- F. Mirror the application code to a third-party Git repository that uses Amazon S3 storage
- G. Use the repository for deployment.
- H. Replace the custom scripts and tools with AWS CodeBuild
- I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 32

- (Exam Topic 1)

A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5.1 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

- A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- B. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
- C. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- D. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains a message
- E. Have the application process each record, and transform the record into JSON format
- F. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
- G. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match
- H. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirement
- I. Define the output format as JSON
- J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
- K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match
- L. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirement
- M. Define the output format as JSON
- N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:

You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object creation events occur. The Lambda function will then trigger the Glue ETL job to transform the records, masking the sensitive data and modifying the output format to

JSON. This solution meets all requirements.

Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON.

Once complete, have the ETL job send the results to another S3 bucket for internal processing. <https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html>
https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipel

NEW QUESTION 36

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to

stream. Each media item Contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enable. The company uses Amazon CloudFront to serve these movie files.

The company has 100.000 media items, and each media item can have many different S3 objects that represent different encodings of the same media S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID

Because of an expiring contract with a media provider, the company must remove 2.000 media Items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours The company must ensure that the content cannot be recovered.

Which combination of actions will meet these requirements? (Select TWO.)

- A. Configure the dynamoDB table with a TTL fiel
- B. Create and invoke an AWS Lambda function to perform a conditional update Set the TTL field to the time of the contract's expiration on every affected media item.
- C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
- D. Write a script to perform a conditional delete on all the affected DynamoDB records
- E. Temporarily suspend versioning on the S3 bucke
- F. Create and invoke an AWS Lambda function that deletes affected objects Reactivate versioning when the operation is complete
- G. Write a script to delete objects from Amazon S3 Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 37

- (Exam Topic 1)

A company hosts a web application that tuns on a group of Amazon EC2 instances that ate behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads lo reverse-engineer a sophisticated attack of the application.

Which approach should the company take to achieve this goal?

- A. Enable VPC Flow Log
- B. Store the flow logs in an Amazon S3 bucket for analysis.
- C. Enable Traffic Mirroring on the network interface of the EC2 instance
- D. Send the mirrored traffic lo a target for storage and analysis.
- E. Create an AWS WAF web AC
- F. and associate it with the AL
- G. Configure AWS WAF logging.
- H. Enable logging for the AL
- I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 39

- (Exam Topic 1)

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.

Which set of actions should a solutions architect take to meet these requirements?

- A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instance
- B. Use Systems Manager to generate patch compliance reports.
- C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
- D. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.
- E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation jo
- F. Use Amazon Inspector to generate patch compliance reports.
- G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
- H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Answer: A

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html>

NEW QUESTION 43

- (Exam Topic 1)

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed. Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organizatio
- B. Use an AWS Systems Manager Parameter Store parameter to store accounnumbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda

function to update the security policy in the Firewall Manager administrative account

C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule

D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resource

E. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.

F. Create AWS WAF rules in the management account of the organization

G. Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs

H. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts

I. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization

J. Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs

K. Create IAM users in member accounts

Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 48

- (Exam Topic 1)

A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.

A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.

What should the solutions architect do next to meet these requirements?

- A. Create the OrganizationAccountAccess IAM group in each member account
- B. Include the necessary IAM roles for each administrator.
- C. Create the OrganizationAccountAccessPolicy IAM policy in each member account
- D. Connect the member accounts to the management account by using cross-account access.
- E. Create the OrganizationAccountAccessRole IAM role in each member account
- F. Grant permission to the management account to assume the IAM role.
- G. Create the OrganizationAccountAccessRole IAM role in the management account
- H. Attach the Administrator Access AWS managed policy to the IAM role
- I. Assign the IAM role to the administrators in each member account.

Answer: C

NEW QUESTION 51

- (Exam Topic 1)

A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes.

Which solution meets the requirements?

- A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behavior)
- B. Send Amazon SNS notifications when anomalous behaviors are detected.
- C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
- D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
- E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda
- F. Create a Lambda function to output records to Amazon Kinesis Data Stream
- G. Analyze any anomalies with Amazon Kinesis Data Analytics
- H. Send SNS notifications when anomalous behaviors are detected.
- I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior
- J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:

[https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection \(it uses AWS proprietary Random Cut Forest Algorithm\)](https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection (it uses AWS proprietary Random Cut Forest Algorithm))

NEW QUESTION 53

- (Exam Topic 1)

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24.

Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.

Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets.

What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0
- D. An outbound rule for port 80 to destination 10.0.0.0/24
- E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

Explanation:

Ephemeral ports are not covered in the syllabus, so be careful that you don't confuse day-to-day best practice with what is required for the exam. Link to an explanation on Ephemeral ports here: <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbCwo4IXefMI7jAnaK/netw>

NEW QUESTION 54

- (Exam Topic 1)

A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability. Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC. and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only. Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A. Create an AWS Transit Gateway
- B. Attach the shared VPC and the authorized business unit VPCs to the transit gateway
- C. Create a single transit gateway route table and associate it with all of the attached VPC
- D. Allow automatic propagation of routes from the attachments into the route table
- E. Configure VPC routing tables to send traffic to the transit gateway.
- F. Create a VPC endpoint service using the centralized application NLB and enable (the option to require endpoint acceptance)
- G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service
- H. Accept authorized endpoint requests from the endpoint service console.
- I. Create a VPC peering connection from each business unit VPC to the shared VPC
- J. Accept the VPC peering connections from the shared VPC console
- K. Configure VPC routing tables to send traffic to the VPC peering connection.
- L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPC
- M. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC
- N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-pre>

NEW QUESTION 58

- (Exam Topic 1)

A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances.

The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result, the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive.

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

- A. Suspend the Auto Scaling group's HealthCheck scaling process
- B. Use Session Manager to log in to an instance that is marked as unhealthy
- C. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.
- D. Set the termination policy to OldestInstance on the Auto Scaling group
- E. Use Session Manager to log in to an instance that is marked as unhealthy
- F. Suspend the Auto Scaling group's Terminate process
- G. Use Session Manager to log in to an instance that is marked as unhealthy

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>

It shows. For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The

Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process!

It is a process along with the following: AddToLoadBalancer, AlarmNotification, AZRebalance, HealthCheck, InstanceRefresh, ReplaceUnhealthy, ScheduledActions.

From the requirements, Some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-r>

NEW QUESTION 59

- (Exam Topic 1)

A scientific organization requires the processing of text and picture data stored in an Amazon S3 bucket. The data is gathered from numerous radar stations during a mission's live, time-critical phase. The data is uploaded by the radar stations to the source S3 bucket. The data is preceded with the identification number of the radar station.

In a second account, the business built a destination S3 bucket. To satisfy a compliance target, data must be transferred from the source S3 bucket to the destination S3 bucket. Replication is accomplished by using an S3 replication rule that covers all items in the source S3 bucket.

A single radar station has been recognized as having the most precise data. At this radar station, data replication must be completed within 30 minutes of the radar station uploading the items to the source S3 bucket.

What actions should a solutions architect take to ensure that these criteria are met?

- A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket
- B. Select to use all available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
- D. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data. Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations. Monitor the maximum replication time to the destination
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold
- F. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint. Monitor the S3 destination bucket's TotalRequestLatency metric. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes

G. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data Enable S3 Replication Time Control (S3 RTC) Monitor the maximum replication time to the destination Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html>

NEW QUESTION 62

- (Exam Topic 1)

A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS for MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.

Which strategy should a solutions architect recommend to remediate these security risks?

- A. Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credential
- B. Take a snapshot of the DB instance and encrypt a copy of that snapshot
- C. Replace the DB instance with a new DB instance that is based on the encrypted snapshot.
- D. Enable IAM DB authentication on the DB instance
- E. Grant the Lambda execution role access to the DB instance
- F. Modify the DB instance and enable encryption.
- G. Enable IAM DB authentication on the DB instance
- H. Grant the Lambda execution role access to the DB instance
- I. Create an encrypted read replica of the DB instance
- J. Promote the encrypted read replica to be the new primary node.
- K. Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameter
- L. Create another Lambda function to automatically rotate the credential
- M. Create an encrypted read replica of the DB instance
- N. Promote the encrypted read replica to be the new primary node.

Answer: A

Explanation:

Parameter store can store DB credentials as secure string but CANNOT rotate secrets, hence, go with A + Cannot enable encryption on existing MySQL RDS instance, must create a new encrypted one from unencrypted snapshot.

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-> Encrypting a unencrypted instance of DB or creating a encrypted replica of an unencrypted DB instance are not possible Hence A is the only solution possible.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption>.

NEW QUESTION 65

- (Exam Topic 1)

A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the front-end instances running behind a load balancing appliance that has a virtual offering on AWS Current*, the operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including.

- A DDoS attack.
- An SQL injection attack
- Several successful dictionary attacks on SSH accounts on the web servers

The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's solutions architects have decided to use the following approach;

- Code review the existing application and fix any SQL injection issues.
- Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching.
- Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.

What additional steps will address all of the identified attack types while providing high availability and minimizing risk?

- A. Enable SSH access to the Amazon EC2 instances using a security group that limits access to specific IP
- B. Migrate on-premises MySQL to Amazon RDS Multi-AZ Install the third-party load balancer from the AWS Marketplace and migrate the existing rules to the load balancer's AWS instances Enable AWS Shield Standard for DDoS protection
- C. Disable SSH access to the Amazon EC2 instance
- D. Migrate on-premises MySQL to Amazon RDS Multi-AZ Leverage an Elastic Load Balancer to spread the load and enable AWS Shield Advanced for protection
- E. Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.
- F. Enable SSH access to the Amazon EC2 instances through a bastion host secured by limiting access to specific IP addresses
- G. Migrate on-premises MySQL to a self-managed EC2 instance
- H. Leverage an AWS Elastic Load Balancer to spread the load, and enable AWS Shield Standard for DDoS protection Add an Amazon CloudFront distribution in front of the website.
- I. Disable SSH access to the EC2 instance
- J. Migrate on-premises MySQL to Amazon RDS Single-A
- K. Leverage an AWS Elastic Load Balancer to spread the load Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules.

Answer: B

NEW QUESTION 69

- (Exam Topic 1)

A solutions architect works for a government agency that has strict disaster recovery requirements All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.

Which solution meets these requirements?

- A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLM) to run once daily to copy the EBS snapshots to the additional Regions.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.

- C. Set up AWS Backup to create the EBS snapshot
- D. Configure Amazon S3 cross-Region replication to copy the EBS snapshots to the additional Regions.
- E. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions.

Answer: B

NEW QUESTION 71

- (Exam Topic 1)

A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.

The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.

What should a solutions architect recommend to meet these requirements?

- A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
- B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- C. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region
- D. Configure the CloudWatch alarm to invoke the Lambda function.
- E. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- F. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy
- G. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
- H. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region
- I. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
- J. Remove the read replica
- K. Replace the read replica with a standalone RDS DB instance
- L. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
- M. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted targets
- N. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
- O. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region
- P. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 74

- (Exam Topic 1)

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.

Which solution will meet these requirements?

- A. Update the Auto Scaling group to use the AZ2 subnet only
- B. Delete and re-create the AZ1 subnet using half the previous address space
- C. Adjust the Auto Scaling group to also use the new AZ1 subnet
- D. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- E. Remove the current AZ2 subnet
- F. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet
- G. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
- H. Terminate the EC2 instances in the AZ1 subnet
- I. Delete and re-create the AZ1 subnet using half the address space
- J. Update the Auto Scaling group to use this new subnet
- K. Repeat this for the second AZ
- L. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.
- M. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ
- N. Update the existing Auto Scaling group to target the new subnets in the new VPC.
- O. Update the Auto Scaling group to use the AZ2 subnet only
- P. Update the AZ1 subnet to have half the previous address space
- Q. Adjust the Auto Scaling group to also use the AZ1 subnet again
- R. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- S. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet
- T. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls

It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

NEW QUESTION 75

- (Exam Topic 1)

A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

- A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.
- B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS QMS.
- C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.
- D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS
- E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

Answer: BD

Explanation:

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/> <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-postgresql-database>

NEW QUESTION 80

- (Exam Topic 1)

A fitness tracking company serves users around the world, with its primary markets in North America and Asia. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:

- Be resilient to problems with the application in any Region.
- Write to a database in a single Region.
- Read from multiple Regions.
- Support resiliency across application tiers in each Region.
- Support the relational database semantics reflected in the application. Which combination of steps should a solutions architect take? (Select TWO.)

- A. Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.
- B. Deploy we
- C. application, and MySQL database servers to Amazon EC2 instances in each Regio
- D. Set up the application so that reads and writes are local to the Regio
- E. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Region
- F. Set upcross-Region replication for the database layer.
- G. Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.
- H. Set up web, application, and Amazon RDS for MySQL instances in each Regio
- I. Set up the application so that reads are local and writes are partitioned based on the use
- J. Set up a Multi-AZ failover for the web, application, and database server
- K. Set up cross-Region replication for the database layer.
- L. Set up active-active web and application servers in each Regio
- M. Deploy an Amazon Aurora global database with clusters in each Regio
- N. Set up the application to use the in-Region Aurora database endpoint
- O. Create snapshots of the web and application servers and store them in an Amazon S3 bucket in both Regions.

Answer: CE

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

Geoproximity routing policy is good to control the user traffic to specific regions. However, a multivalue answer routing policy may cause the users to be randomly sent to other healthy regions that may be far away from the user's location. You can use geolocation routing policy to direct the North American users to your servers on the North America region and configure failover routing to the Asia region in case the North America region fails. You can configure the same for the Asian users pointed to the Asia region servers and have the North America region as its backup.

NEW QUESTION 84

- (Exam Topic 1)

An online retail company hosts its stateful web-based application and MySQL database in an on-premises data center on a single server. The company wants to increase its customer base by conducting more marketing campaigns and promotions. In preparation, the company wants to migrate its application and database to AWS to increase the reliability of its architecture.

Which solution should provide the HIGHEST level of reliability?

- A. Migrate the database to an Amazon RDS MySQL Multi-AZ DB instanc
- B. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
- C. Store sessions in Amazon Neptune.
- D. Migrate the database to Amazon Aurora MySQ
- E. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
- F. Store sessions in an Amazon ElastiCache for Redis replication group.
- G. Migrate the database to Amazon DocumentDB (with MongoDB compatibility). Deploy the application in an Auto Scaling group on Amazon EC2 instances behind a Network Load Balance
- H. Store sessions in Amazon Kinesis Data Firehose.
- I. Migrate the database to an Amazon RDS MariaDB Multi-AZ DB instanc
- J. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
- K. Store sessions in Amazon ElastiCache for Memcached.

Answer: B

NEW QUESTION 87

- (Exam Topic 1)

To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet. How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

- A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data.
- B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- C. Use inter-region VPC peering to access the data in other AWS Regions.
- D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- E. Use an AWS transit VPC solution to access data in other AWS Regions.
- F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:

This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. <https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>
<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g>

NEW QUESTION 88

- (Exam Topic 1)

A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware and Hyper-V. An administrator needs to select the correct services to collect data (or the initial migration discovery process). The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data.

Which solution meets these requirements?

- A. Use the AWS Agentless Discovery Connector for data collection on physical servers and all VMs.
- B. Store the collected data in Amazon S3. Query the data with S3 Select.
- C. Generate reports by using Kibana hosted on Amazon EC2.
- D. Use the AWS Application Discovery Service agent for data collection on physical servers and all VMs. Store the collected data in Amazon Elastic File System (Amazon EFS). Query the data and generate reports with Amazon Athena.
- E. Use the AWS Application Discovery Service agent for data collection on physical servers and Hyper-V.
- F. Use the AWS Agentless Discovery Connector for data collection on VMware.
- G. Store the collected data in Amazon S3. Query the data with Amazon Athena.
- H. Generate reports by using Amazon QuickSight.
- I. Use the AWS Systems Manager agent for data collection on physical server.
- J. Use the AWS Agentless Discovery Connector for data collection on all VMs.
- K. Store, query, and generate reports from the collected data by using Amazon Redshift.

Answer: C

Explanation:

<https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html> <https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-connector.html>

NEW QUESTION 90

- (Exam Topic 1)

A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.

The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.

Which solution meets these requirements?

- A. Create an AWS PrivateLink interface VPC endpoint.
- B. Connect this endpoint to the endpoint service that the third-party SaaS application provides.
- C. Create a security group to limit the access to the endpoint.
- D. Associate the security group with the endpoint.
- E. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VPC.
- F. Configure network ACLs to limit access across the VPN tunnels.
- G. Create a VPC peering connection between the third-party SaaS application and the company VPC. Update route tables by adding the needed routes for the peering connection.
- H. Create an AWS PrivateLink endpoint service.
- I. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint service.
- J. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Answer: A

Explanation:

Reference architecture - <https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html> Note from documentation that Interface Endpoint is at client side.

NEW QUESTION 93

- (Exam Topic 1)

A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.

After initial deployment, the company observes 1,000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.

Which approach should the company take to secure its API?

- A. Create an Amazon CloudFront distribution with the API as the origin.
- B. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day.
- C. Associate the web ACL with the CloudFront distribution.

- D. Configure CloudFront with an origin access identity (OAI) and associate it with the distributio
- E. Configure API Gateway to ensure only the OAI can execute the POST method.
- F. Create an Amazon CloudFront distribution with the API as the origi
- G. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per da
- H. Associate the web ACL with the CloudFront distributio
- I. Add a custom header to the CloudFront distribution populated with an API ke
- J. Configure the API to require an API key on the POST method.
- K. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the AP
- L. Create a resource policy with a request limit and associate it with the AP
- M. Configure the API to require an API key on the POST method.
- N. Associate the web ACL with the AP
- O. Create a usage plan with a request limit and associate it with the AP
- P. Create an API key and add it to the usage plan.

Answer: D

Explanation:

"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys."

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

NEW QUESTION 98

- (Exam Topic 1)

A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.

Which approach meets these requirements with the LEAST amount of operational overhead?

- A. Configure AWS Systems Manager OpsCenter to ingest operational events from the on-premises systems Retire the on-premises management application and adopt OpsCenter as the hub
- B. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Personal Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to the HTTPS endpoint of the management application
- C. Modify the on-premises management application to call the AWS Health API to poll for status events of AWS services.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Service Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to an HTTPS endpoint for the management application with a topic filter corresponding to the services being used

Answer: A

Explanation:

ALB & NLB both supports IPs as targets. Questions is based on TCP traffic over VPN to on-premise. TCP is layer 4 and the , load balancer should be NLB. Then next questions does NLB supports loadbalcning traffic over VPN. And answer is YEs based on below URL.

<https://aws.amazon.com/about-aws/whats-new/2018/09/network-load-balancer-now-supports-aws-vpn/>

Target as IPs for NLB & ALB: <https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5> <https://aws.amazon.com/elasticloadbalancing/application-load-balancer/>

NEW QUESTION 101

- (Exam Topic 2)

A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.

The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.

What could be the cause of the error messages for these customers?

- A. The Lambda function reached its concurrency limit.
- B. The Lambda function its Region limit for concurrency.
- C. The company reached its API Gateway account limit for calls per second.
- D. The company reached its API Gateway default per-method limit for calls per second.

Answer: C

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-reques>

NEW QUESTION 102

- (Exam Topic 2)

A company wants to improve cost awareness for its Amazon EMR platform The company has aWocated budgets for each team's Amazon EMR usage When a budgetary threshold is reached a notification should be sent by email to the budget office's distribution list Teams should be able lo view their EMR cluster expenses to date A solutions architect needs to create a solution that ensures this policy is proactively and centrally enforced in a multi-account environment Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

- A. Update the AWS CloudFormation template to include the AWS Budgets Budget resource with the NotificationsWithSubscnbers property
- B. Implement Amazon CloudWatch dashboards for Amazon EMR usage
- C. Create an EMR bootstrap action that runs at startup that calls the Cost Explorer API to set the budget on the cluster with the GetCostForecast and NotificationsWithSubscnbers actions
- D. Create an AWS Service Catalog portfolio for each tea
- E. Add each team's Amazon EMR cluster as an AWS CloudFormation template to their Service Catalog portfolio as a Product

F. Create an Amazon CloudWatch metric for billing Create a custom alert when costs exceed the budgetary threshold.

Answer: BE

NEW QUESTION 106

- (Exam Topic 2)

A company plans to refactor a monolithic application into a modern application designed deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements

- It should allow changes to be released several times every hour.
- * It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations Deploy the application by replacing Amazon EC2 instances
- B. Specify AWS Elastic Beanstalk to stage in a secondary environment as the deployment target for the CI/CD pipeline of the application
- C. To deploy swap the staging and production environment URLs.
- D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
- E. Roll out application updates as part of an Auto Scaling event using prebuilt AMI
- F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:

It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 107

- (Exam Topic 2)

A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs ETL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.

The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data. The customers also need access to the most recent data when the company publishes the data. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Data Exchange for APIs to share data with customer
- B. Configure subscription verification In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshift
- C. Require the data customers to subscribe to the data product In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift
- D. cluster
- E. Configure subscription verification
- F. Require the data customers to subscribe to the data product.
- G. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically
- H. Use AWS Data Exchange for S3 to share data with customers.
- I. Configure subscription verification
- J. Require the data customers to subscribe to the data product Publish the Amazon Redshift data to an Open Data on AWS Data Exchange
- K. Require the customers to subscribe to the data product in AWS Data Exchange
- L. In the AWS account of the company that produces the data, attach IAM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.

Answer: D

NEW QUESTION 111

- (Exam Topic 2)

A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.

A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances
- B. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.
- C. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances Ensure that the EC2 instances are configured in unlimited mode.
- D. Modify the DB instance to create a read replica in the same Availability Zone
- E. Promote the read replica to be the primary DB instance in failure scenarios.
- F. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.
- G. Create a replication group for the ElastiCache for Redis cluster
- H. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.
- I. Create a replication group for the ElastiCache for Redis cluster
- J. Enable Multi-AZ on the cluster.

Answer: ADE

NEW QUESTION 116

- (Exam Topic 2)

A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.

The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DownloadUpload",
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::BucketName/*"
    },
    {
      "Sid": "KMSAccess",
      "Action": [
        "kms:Decrypt",
        "kms:Encrypt"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:kms:Region:Account:key/Key ID"
    }
  ]
}
```

During tests, the solutions architect was able to successfully get existing test objects in the S3 bucket. However, attempts to upload a new object resulted in an error message. The error message stated that the action was forbidden.

Which action must the solutions architect add to the IAM policy to meet all the requirements?

- A. kms:GenerateDataKey
- B. kms:GetKeyPolicy
- C. kms:GetPublicKey
- D. kms:SKJN

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-access-denied-error-kms/>

"An error occurred (AccessDenied) when calling the PutObject operation: Access Denied" This error message indicates that your IAM user or role needs permission for the kms:GenerateDataKey action.

NEW QUESTION 117

- (Exam Topic 2)

A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows:

GET /posts/{postId} to get post details
GET /users/{userId} to get user details

GET /comments/{commentId} to get comments details

The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by making the comments appear in real time.

Which design should be used to reduce comment latency and improve user experience?

- A. Use edge-optimized API with Amazon CloudFront to cache API responses.
- B. Modify the blog application code to request GET/commentsV{commentId} every 10 seconds.
- C. Use AWS AppSync and leverage WebSockets to deliver comments.
- D. Change the concurrency limit of the Lambda functions to lower the API response time.

Answer: C

NEW QUESTION 121

- (Exam Topic 2)

A finance company is storing financial records in an Amazon S3 bucket. The company persists a record for every financial transaction. According to regulatory requirements, the records cannot be modified for at least 1 year after they are written. The records are read on a regular basis and must be immediately accessible.

Which solution will meet these requirements?

- A. Create a new S3 bucket.
- B. Turn on S3 Object Lock, set a default retention period of 1 year, and set the retention mode to compliance mode.
- C. Store all records in the new S3 bucket.
- D. Create an S3 Lifecycle rule to immediately transfer new objects to the S3 Glacier storage tier. Create an S3 Glacier Vault Lock policy that has a retention period of 1 year.
- E. Create an S3 Lifecycle rule to immediately transfer new objects to the S3 Intelligent-Tiering storage tier. Set a retention period of 1 year.
- F. Create an S3 bucket policy with a Deny action for PutObject operations with a condition where the s3:x-amz-object-retention header is not equal to 1 year.

Answer: A

NEW QUESTION 122

- (Exam Topic 2)

A large education company recently introduced Amazon Workspaces to provide access to internal applications across multiple universities. The company is storing user proxies on an Amazon FSx for Windows File Server file system. The file system is configured with a DNS alias and is connected to a self-managed Active Directory. As more users begin to use the Workspaces login time increases to unacceptable levels.

An investigation reveals a degradation in performance of the file system. The company created the file system on HDD storage with a throughput of 16 MBps. A solutions architect must improve the performance of the file system during a defined maintenance window.

What should the solutions architect do to meet these requirements with the LEAST administrative effort?

- A. Use AWS Backup to create a point-in-time backup of the file system. Restore the backup to a new FSx for Windows File Server file system. Select SSD as the storage type. Select 32 MBps as the throughput capacity. When the backup and restore process is completed, adjust the DNS alias accordingly. Delete the original file system.
- B. Disconnect users from the file system. In the Amazon FSx console, update the throughput capacity to 32 MBps. Update the storage type to SSD. Reconnect users to the file system.
- C. Deploy an AWS DataSync agent onto a new Amazon EC2 instance.
- D. Create a task. Configure the existing file system as the source location. Configure a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput as the target location. Schedule the task. When the task is completed, adjust the DNS alias accordingly. Delete the original file system.
- E. Enable shadow copies on the existing file system by using a Windows PowerShell command. Schedule the shadow copy job to create a point-in-time backup of the file system. Choose to restore previous versions. Create a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput. When the copy job is completed, adjust the DNS alias. Delete the original file system.

Answer: D

NEW QUESTION 123

- (Exam Topic 2)

A solutions architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month. Which combination of steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Select THREE.)

- A. Implement an IAM policy that requires users to specify a 'workload' tag for cost allocation when launching Amazon EC2 instances.
- B. Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types.
- C. Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage.
- D. Place conditions in the users' IAM policies that limit the number of instances they are able to launch.
- E. Define 'workload' as a cost allocation tag in the AWS Billing and Cost Management console.
- F. Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost.

Answer: AEF

NEW QUESTION 125

- (Exam Topic 2)

A company manages hundreds of AWS accounts centrally in an organization using AWS Organizations. The company recently started to allow product teams to create and manage their own S3 access points in their accounts. The S3 access points can be accessed only within VPCs, not on the internet.

What is the MOST operationally efficient way to enforce this requirement?

- A. Set the S3 access point resource policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- B. Create an SCP at the root level in the organization to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- C. Use AWS CloudFormation StackSets to create a new IAM policy in each AWS account that allows the s3:CreateAccessPoint action only if the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
- D. Set the S3 bucket policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

Answer: A

NEW QUESTION 127

- (Exam Topic 2)

A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.

The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.

Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.

How can the company solve this problem?

- A. Turn on termination protection for the EC2 instances.
- B. Update the visibility timeout for the SQS queue to 3 hours.
- C. Configure scale-in protection for the instances during processing.
- D. Update the redrive policy and set maxReceiveCount to 0.

Answer: D

NEW QUESTION 130

- (Exam Topic 2)

A company has an application. Once a month, the application creates a compressed file that contains every object within an Amazon S3 bucket. The total size of the objects before compression is 1 TB.

The application runs by using a scheduled cron job on an Amazon EC2 instance that has a 5 TB Amazon Elastic Block Store (Amazon EBS) volume attached. The application downloads all the files from the source S3 bucket to the EBS volume, compresses the file, and uploads the file to a target S3 bucket. Every invocation of the application takes 2 hours from start to finish.

Which combination of actions should a solutions architect take to OPTIMIZE costs for this application? (Select TWO.)

- A. Migrate the application to run an AWS Lambda function Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the Lambda function to run once each month
- B. Configure the application to download the source files by using streams Direct the streams into a compression library Direct the output of the compression library into a target object in Amazon S3
- C. Configure the application to download the source files from Amazon S3 and save the files to local storage Compress the files and upload them to Amazon S3
- D. Configure the application to run as a container in AWS Fargate Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the task to run once each month
- E. Provision an Amazon Elastic File System (Amazon EFS) file system Attach the file system to the AWS Lambda function

Answer: CD

NEW QUESTION 134

- (Exam Topic 2)

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

What should the solutions architect do to meet this requirement with the LEAST amount of effort?

- A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API call
- B. Create an inventory of the required API calls and resources for each Lambda function
- C. Create new IAM access policies for each Lambda function
- D. Review the new policies to ensure that they meet the company's business requirements.
- E. Turn on AWS CloudTrail logging for the AWS account
- F. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail log. Review the generated policies to ensure that they meet the company's business requirements.
- G. Turn on AWS CloudTrail logging for the AWS account
- H. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report
- I. Review the report
- J. Create IAM access policies that provide more restrictive permissions for each Lambda function.
- K. Turn on AWS CloudTrail logging for the AWS account
- L. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each execution role
- M. Create a new IAM access policy for each role
- N. Export the generated roles to an S3 bucket
- O. Review the generated policies to ensure that they meet the company's business requirements.

Answer: B

Explanation:

IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

NEW QUESTION 137

- (Exam Topic 2)

A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.

The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sales team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new S3 bucket in the marketing account
- B. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing account
- C. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.
- D. Create an SCP to grant access to the S3 bucket to the marketing account
- E. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sales account with the marketing account
- F. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
- G. Update the S3 bucket policy in the marketing account to grant access to the QuickSight role
- H. Create a KMS grant for the encryption key that is used in the S3 bucket
- I. Grant decrypt access to the QuickSight role
- J. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
- K. Create an IAM role in the sales account and grant access to the S3 bucket
- L. From the marketing account, assume the IAM role in the sales account to access the S3 bucket
- M. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.

Answer: D

NEW QUESTION 138

- (Exam Topic 2)

A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.

Which solution will meet these requirements?

- A. Migrate the database to Amazon DynamoD
- B. Store the leader different table
- C. Use Apache HiveQLJOIN statements to build the leaderboard
- D. Keep the leaderboard data in the RDS DB instanc
- E. Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.
- F. Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destinatio
- G. Query the S3 bucket by using Amazon Athena for the leaderboard.
- H. Add a read-only replica to the RDS DB instanc
- I. Add an RDS Proxy database proxy.

Answer: C

NEW QUESTION 142

- (Exam Topic 2)

A solutions architect needs to provide AWS Cost and Usage Report data from a company's AWS Organizations management account. The company already has an Amazon S3 bucket to store the reports. The reports must be automatically ingested into a database that can be visualized with other tools.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that a new object creation in the S3 bucket will trigger
- B. Create an AWS Cost and Usage Report configuration to deliver the data into the S3 bucket
- C. Configure an AWS Glue crawler that a new object creation in the S3 bucket will trigger.
- D. Create an AWS Lambda function that a new object creation in the S3 bucket will trigger
- E. Create an AWS Glue crawler that the AWS Lambda function will trigger to crawl objects in the S3 bucket
- F. Create an AWS Glue crawler that the Amazon EventBridge (Amazon CloudWatch Events) rule will trigger to crawl objects in the S3 bucket

Answer: BDF

NEW QUESTION 145

- (Exam Topic 2)

A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, master, and core nodes. The EMR tasks run each morning, starting at 1:00 AM, and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.

The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

- A. Launch all task, master, and core nodes on Spot Instances in an instance fleet
- B. Terminate the cluster, including all instances, when the processing is completed.
- C. Launch the master and core nodes on On-Demand Instance
- D. Launch the task nodes on Spot Instances in an instance fleet
- E. Terminate the cluster, including all instances, when the processing is complete
- F. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- G. Continue to launch all nodes on On-Demand Instance
- H. Terminate the cluster
- I. Terminate all instances, when the processing is complete
- J. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- K. Launch the master and core nodes on On-Demand Instance
- L. Launch the task nodes on Spot Instances in an instance fleet
- M. Terminate only the task node Instances when the processing is completed. Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: B

NEW QUESTION 147

- (Exam Topic 2)

A company wants to migrate its workloads from on-premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises infrastructure that consists of physical machines and VMs that host numerous applications.

The company must capture details about the system configuration, system performance, running processes, and network connections of its on-premises servers. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.

Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

- A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
- B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs
- C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
- D. Group servers into applications for migration by using AWS Migration Hub.
- E. Generate recommended instance types and associated costs by using AWS Migration Hub.
- F. Import data about server sizes into AWS Trusted Advisor
- G. Follow the recommendations for cost optimization.

Answer: BDF

NEW QUESTION 152

- (Exam Topic 2)

A company has automated the nightly retraining of its machine learning models by using AWS Step Functions. The workflow consists of multiple steps that use AWS Lambda. Each step can fail for various reasons, and any failure causes a failure of the overall workflow.

A review reveals that the retraining has failed multiple nights in a row without the company noticing the failure. A solutions architect needs to improve the workflow so that notifications are sent for all types of failures in the retraining process.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list.
- B. Create a task named "Email" that forwards the input arguments to the SNS topic
- C. Add a Catch field to all Tasks
- D. Map
- E. and Parallel states that have a statement of "ErrorEquals": ["states.all"] and "Next": "Email".
- F. Add a new email address to Amazon Simple Email Service (Amazon SES). Verify the email address.
- G. Create a task named "Email" that forwards the input arguments to the SES email address
- H. Add a Catch field to all Task, Map, and Parallel states that have a statement of "ErrorEquals": ["states.Bun time"] and "Next": "Email".

Answer: BCD

NEW QUESTION 153

- (Exam Topic 2)

A company wants to migrate its data analytics environment from on premises to AWS. The environment consists of two simple Node.js applications. One of the applications collects sensor data and loads it into a MySQL database. The other application aggregates the data into reports. When the aggregation jobs run, some of the load jobs fail to run correctly.

The company must resolve the data loading issue. The company also needs the migration to occur without interruptions or changes for the company's customers. What should a solutions architect do to meet these requirements?

- A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind a Network Load Balancer (NLB). and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance.
- B. Point the collector DNS record to the NLB.
- C. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Move the aggregation jobs to run against the Aurora MySQL database. Set up collection endpoints behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
- D. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB) and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
- E. Set up an Amazon Aurora MySQL database. Create an Aurora Replica for the Aurora MySQL database and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as an Amazon Kinesis data stream. Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the Kinesis data stream.

Answer: C

Explanation:

Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%.

NEW QUESTION 156

- (Exam Topic 2)

A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.

The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network. Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.

The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts. Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute changes by publishing messages to its SNS topic.
- B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.
- C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR range.
- D. Share the customer-managed prefix list.... organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.
- E. Create an IAM role in each account in the organization.
- F. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account.
- G. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Answer: C

NEW QUESTION 157

- (Exam Topic 2)

A company has multiple business units. Each business unit has its own AWS account and runs a single website within that account. The company also has a single logging account. Logs from each business unit website are aggregated into a single Amazon S3 bucket in the logging account. The S3 bucket policy provides each business unit with access to write data into the bucket and requires data to be encrypted.

The company needs to encrypt logs uploaded into the bucket using a Single AWS Key Management Service (AWS KMS) CMK. The CMK that protects the data must be rotated once every 365 days.

Which strategy is the MOST operationally efficient for the company to use to meet these requirements?

- A. Create a customer managed CMK in the logging account Update the CMK key policy to provide access to the logging account only Manually rotate the CMK every 365 days.
- B. Create a customer managed CMK in the logging account
- C. Update the CMK key policy to provide access to the logging account and business unit account
- D. Enable automatic rotation of the CMK
- E. Use an AWS managed CMK in the logging account
- F. Update the CMK key policy to provide access to the logging account and business unit accounts Manually rotate the CMK every 365 days.
- G. Use an AWS managed CMK in the logging account Update the CMK key policy to provide access to the logging account only
- H. Enable automatic rotation of the CMK.

Answer: A

NEW QUESTION 159

- (Exam Topic 2)

A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.

Users who access the app from Australia have experienced uploads that take long periods of time Sometimes the files fail to completely upload for these users . A solutions architect must improve the app's performance for these uploads

Which solutions will meet these requirements? (Select TWO.)

- A. Enable S3 Transfer Acceleration on the S3 bucket Configure the app to use the Transfer Acceleration endpoint for uploads
- B. Configure an S3 bucket in each Region to receive the upload
- C. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
- D. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
- E. Configure the app to break the video files into chunks Use a multipart upload to transfer files to Amazon S3.
- F. Modify the app to add random prefixes to the files before uploading

Answer: AD

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/>

NEW QUESTION 163

- (Exam Topic 2)

A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.

The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.

Which solution will meet these requirements?

- A. Create a private VIF from the DX-A connection into a Direct Connect gateway
- B. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availability
- C. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gateway
- D. Peer the transit gateways with each other to support cross-Region routing.
- E. Create a transit VIF from the DX-A connection into a Direct Connect gateway
- F. Associate the eu-west-1 transit gateway with this Direct Connect gateway
- G. Create a transit VIF from the DX-B connection into a separate Direct Connect gateway
- H. Associate the us-east-1 transit gateway with this separate Direct Connect gateway
- I. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
- J. Create a transit VIF from the DX-A connection into a Direct Connect gateway
- K. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability
- L. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway
- M. Configure the Direct Connect gateway to route traffic between the transit gateways.
- N. Create a transit VIF from the DX-A connection into a Direct Connect gateway
- O. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability
- P. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway
- Q. Peer the transit gateways with each other to support cross-Region routing.

Answer: D

NEW QUESTION 168

- (Exam Topic 2)

A greeting card company recently advertised that customers could send cards to their favourite celebrities through the company's platform Since the advertisement was published, the platform has received constant traffic from 10,000 unique users each second.

The platform runs on m5.xlarge Amazon EC2 instances behind an Application Load Balancer (ALB) The instances run in an Auto Scaling group and use a custom AMI that is based on Amazon Linux. The platform uses a highly available Amazon Aurora MySQL DB cluster that uses primary and reader endpoints The platform also uses an Amazon ElastiCache for Redis cluster that uses its cluster endpoint

The platform generates a new process for each customer and holds open database connections to MySQL for the duration of each customer's session However, resource usage for the platform is low.

Many customers are reporting errors when they connect to the platform Logs show that connections to the Aurora database are failing Amazon CloudWatch metrics show that the CPU load is low across the platform and that connections to the platform are successful through the ALB.

Which solution will remediate the errors MOST cost-effectively?

- A. Set up an Amazon CloudFront distribution Set the ALB as the origin Move all customer traffic to the CloudFront distribution endpoint
- B. Use Amazon RDS Proxy Reconfigure the database connections to use the proxy
- C. Increase the number of reader nodes in the Aurora MySQL cluster

D. Increase the number of nodes in the ElastiCache for Redis cluster

Answer: C

NEW QUESTION 170

- (Exam Topic 2)

A company hosts a web application on AWS in the us-east-1 Region. The application servers are distributed across three Availability Zones behind an Application Load Balancer. The database is hosted in a MySQL database on an Amazon EC2 instance. A solutions architect needs to design a cross-Region data recovery solution using AWS services with an RTO of less than 5 minutes and an RPO of less than 1 minute. The solutions architect is deploying application servers in us-west-2 and has configured Amazon Route 53 health checks and DNS failover to us-west-2.

Which additional step should the solutions architect take?

- A. Migrate the database to an Amazon RDS for MySQL instance with a cross-Region read replica in us-west-2
- B. Migrate the database to an Amazon Aurora global database with the primary in us-east-1 and the secondary in us-west-2
- C. Migrate the database to an Amazon RDS for MySQL instance with a Multi-AZ deployment
- D. Create a MySQL standby database on an Amazon EC2 instance in us-west-2

Answer: C

NEW QUESTION 174

.....

Thank You for Trying Our Product

* **100% Pass or Money Back**

All our products come with a 90-day Money Back Guarantee.

* **One year free update**

You can enjoy free update one year. 24x7 online support.

* **Trusted by Millions**

We currently serve more than 30,000,000 customers.

* **Shop Securely**

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Certified-Solutions-Architect-Professional-dumps.html>